THE VALUE PROPOSITION FOR FRACTIONATED SPACE ARCHITECTURES*

Owen Brown** Defense Advanced Research Projects Agency

Paul Eremenko† Booz Allen Hamilton

THE PROBLEM

The complexity of a spacecraft, as with any other engineering system, is driven by the twin objectives of delivering a particular capability and of doing so robustly in the face of uncertainty. Whereas the intrinsic difficulty of the underlying mission minimally bounds the complexity of the engineering system needed to effect it, in reality such a minimalist system would be of little practical utility; it is not enough to deliver a given capability – it must be delivered with some degree of robustness in the face of various sources of risk or uncertainty.¹

In the particular case of space systems, the array of such sources of uncertainty is both vast and diverse. They include, for instance, technical uncertainties encompassing the risk of component failure, a software bug, a design flaw, a launch vehicle failure, or an erroneous command (if the operator is loosely construed to be part of the spacecraft system), as well as environmental uncertainties such as variations beyond some nominal range in the environmental conditions during spacecraft operations, including temperature, radiation levels, space object impact, etc.

Additionally, there are programmatic sources of uncertainty to which a successful system must also exhibit robustness. One example is the demand for the capability or service provided by the spacecraft during its operational life. Demand fluctuations can occur due to a variety of factors including a change in user constituency, competing providers of the same service, or obsolescence. Another example is requirements uncertainty which, throughout the development of a spacecraft, can necessitate design changes with associated cost,

^{*} This paper is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Approved for public release. Distribution unlimited.

^{**} Program Manager, Tactical Technology Office.

[†] Associate, Defense Business Segment. Corresponding author: Booz Allen Hamilton Inc., 3811 North Fairfax Drive, Suite 600, Arlington, Virginia 22203-1707, tel: (703) 816-5492, fax: (703) 816-5444, e-mail: eremenko@alum.mit.edu.

¹ We use the terms risk and uncertainty interchangeably. Our working definition for risk is variability in an environmental factor (which can be either exogenous or endogenous to the system) that induces volatility in some measure of the system's cost, schedule, or value (which, as discussed *infra*, encompasses capability).

schedule, and value penalties. And of course there is usually uncertainty in the available funding for the development of a particular system. The funding stream, therefore, can fluctuate in a quasi-random manner due to changes in political support, alternative priorities, public perception of the program, and innumerable other factors. Finally, there is fragility. Fragility is not *per se* a source of uncertainty; it is instead an emergent characteristic of complex systems.² Fragility is the tendency of complex systems to exhibit unmodeled failure modes, usually due to an unanticipated component interaction leading to a catastrophic, albeit improbable, sequence of events.

The current approach - or lack thereof - to designing space systems for robustness to uncertainty is the key to the sharply escalating costs and development timelines facing the space industry. Consider, for instance, the typical status quo approach to making a spacecraft robust to technical or environmental risk. The solution is adding margins or parallel redundant component strands in mission-critical areas. This addition of components, in turn, increases system complexity. As complexity grows, so does the system size, cost, and schedule. As a self-fulfilling consequence of the system's increased cost, a higher degree of reliability tends to be imposed as an exogenous, if not occasionally arbitrary, requirement to insure against a catastrophic loss. This results in additional margins and redundancy, and also leads commercial operators to increase costs further through the purchase launch insurance, with costs running from 10% to 20% of the replacement cost of the payload, and onorbit insurance in the range of 2% to 5% annually.3 Military missions, although self-insured in an economic sense, frequently embark on the ultimate exercise in rendundancy and maintain a fully operable "spare" for just about twice the cost of the original spacecraft.

And so commences the cost-complexity death spiral. For while margins, redundancy, and increased reliability requirements are the customary – if inadequate – means for addressing technical and environmental risks, no systematic means at all are used employed to mitigate programmatic risks. So, for instance, a change in requirements or available funding can result in near-total redesign, since most spacecraft are deeply integrated and their performance not easily scalable. The ultimate system cost is therefore many times the initial estimate.

² Carlson, J.M. & Doyle, J., "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 1, pp. 2538-2545 (2002).

³ Anonymous, "Commercial Space and Launch Insurance: Current Market and Future Outlook," *Commercial Space Transportation Quarterly Launch Report*, Federal Aviation Administration, pp. 8-15 (Q4 2002).

Whereas in principle the complexity and the cost of an engineering system should scale roughly in proportion to the system's capability, in practice this is almost never the case. The assured delivery of the capability necessitates making the system robust to various uncertainties. The array of uncertainties and failure modes itself grows with the system's complexity, and the mechanisms for addressing these potential failure modes add to it, with the resultant effect of making overall system complexity grow exponentially. The system's cost follows suit. This is what we term the cost-complexity death spiral.

This tendency has been exacerbated by a requirement-centric minimum-cost acquisition paradigm utilized by the aerospace industry's star client – the U.S. government, and the adoption of similar practices in the commercial aerospace sector – perhaps due to organizational inertia or aversion to bearing the risk which might be associated with process changes. This paradigm has led system architects and decisionmakers to reach the erroneous conclusion that the answer to cost growth is greater capability.⁴ Consequently, longer lifetimes, more transponders per satellite, and multi-functional payloads have all been touted as panacea for the rising cost problem. Unsurprisingly, longer lifetimes have levied additional robustness requirements in the face of system obsolescence, additional capability has led to bigger satellites posing commensurately harder integration, testing, and launch problems, and multi-functional payloads have levied the most stringent of payload pointing and isolation environments across the entire system. The result has been further cost growth.

Perhaps an even more sinister byproduct of rising system complexity is fragility. As we noted *supra*, fragility is the tendency of complex systems to exhibit "emergent" – i.e., unmodeled – failure modes, usually due to an unanticipated component interaction leading to a highly improbable but catastrophic sequence of events. Whereas a complex system can be made robust by anticipating uncertainty and designing for it, fragility tends to rear its ugly head in the most robust, scrupulously designed, and meticulously tested of systems. One need only look to the *Apollo 13*, *Challenger*, or *Columbia* accidents for examples. It is currently addressed through increasingly more expansive simulation and testing. However this approach does not address the root cause of fragility, but attempts – with limited success – to diagnose and root out specific instances of it.

Is the cost-complexity death spiral in spacecraft avoidable? Are uncertainty and fragility manageable phenomena? Can satellites be designed "ground up" for robustness and responsiveness to uncertainty? We believe the answer is answer to all of these questions is a resounding "yes." Prerequisite, however, is an

⁴ See, e.g., Saleh, J., "Flawed metrics*: satellite cost per transponder and cost per operational day (*for guiding design decisions)," *IEEE Transactions on Aerospace and Electronic Systems* (accepted), Vol. --, No. -- (2006).

expansion of the mindset and toolset used to address risk at every stage of the spacecraft lifecycle.

We propose one such tool. It involves the metamorphosis of the gargantuan monolithic spacecraft of today into distributed networks of small fractionated spacecraft components.⁵ This architectural transformation is likely to prove painful for an aerospace industry that has grown remarkably risk averse, but perhaps nowhere near as painful as the paradigmatic shift in acquisition strategy which must accompany it. For in order for such fractionated architectures to become palpable and to be optimized, they must be designed and procured with a value-centric mindset – which lies in stark contrast to the requirement-centric minimum-cost approach utilized today. But before we discuss the implementation – i.e., the value-centric acquisition framework and the technological enablers of fractionated spacecraft – we turn our attention to the particularized value discriminators of fractionated systems, and also briefly address their cost implications.

THE SOLUTION

The Concept of Fractionation

We use "fractionation" as a term of art to describe the decomposition of a system – here a spacecraft – into modules which interact wirelessly to deliver the capability of the original monolithic system.⁶ One can envision the fractionation trade space to be defined by three high-level metrics. First, the heterogeneous degree of fractionation is the number of functionally dissimilar modules into which a system is decomposed. Thus, for instance, a spacecraft with a separate payload, telemetry and communications (T&C), and computation and data handling (C&DH) modules would be fractionated into three heterogeneous modules. Second, the homogeneous degree of fractionation reflects the number of identical modules of a particular type. One could envision a spacecraft whose effective capability would be delivered by a handful of smaller, but otherwise similarly functional modules; or more interestingly – a heterogeneously

⁵ For a more thorough discussion of the various types of uncertainties encountered by space systems and how the behavior of fractionated systems under uncertainty differs from that of monolithic systems, see Brown, O. & Eremenko, P., "Fractionated Space Architectures: A Vision for Response Space," AIAA-RS4-2006-1002, *AIAA* 4th Responsive Space Conference, Los Angeles, CA (2006).

⁶ We are careful to distinguish our notion of fractionation from the more commonly seen concept of modular spacecraft, whereby spacecraft are assembled into a rigid whole through standardized "Lego-like" connections. We believe that the technical requirements and associated risks of orbital rendezvous and docking intrinsic to such a concept, combined with the largely undesirable effects of transmitting forces and torques between various spacecraft modules favor fractionation over modularity.

fractionated spacecraft as described in the example above, with multiple homogeneous modules of a particular type, e.g., C&DH. And the third top-level dimension of fractionation is the type of connectivity among the modules.⁷ The modules could be connected by data links, for instance. Or they could also remotely determine and exchange attitude and position information. Likewise, they could transmit power among themselves, or even remotely effect forces and torques.⁸

Why would fractionation ever be desirable? At first blush, it appears that a fractionated system is hopelessly more complex, heavier, and likely costlier than its monolithic counterpart – necessitating a multitude of transceivers, interfaces, and duplicate components. We address these architectural issues *infra*, but first we turn to some of the advantages that may accrue consequent to fractionation.

Real Options - Design for Architectural Flexibility

Military and commercial satellite operators alike rely on user demand projections to size the capability of a spacecraft. What happens if the demand projections are wrong? If demand is underestimated, the spacecraft will operate at full capacity, and a user prioritization scheme – either through market pricing or through explicit precedence ordering – would have to be implemented. The operator could also fabricate and launch a second spacecraft – assuming demand was underestimated by nearly a factor of two – and hope that by the time fabrication is complete and his turn comes up in the launch queue the demand persists. But if the projection error is smaller, the commercial operator will forego part of its market share to competition or price a certain segment of the user community out of the service, while the military operator will leave a certain user segment un- or under-served. This is the better of the two scenarios. Overestimation of future demand for satellite services has been known to cause business catastrophes exemplified by the bankruptcy of Iridium (and subsequently

⁷ Obviously we do not mean to suggest that these three "metrics" are truly independent, orthogonal axes spanning a trade space in its rigorous sense. For instance the type of intermodule connectivity would be closely related to the heterogeneous degree of fractionation since power exchange would presumably imply a separate power module (or multitude thereof), etc.

⁸ It may be instructive to entertain the extrema of fractionation. One can imagine a spacecraft fractionated into microscopic components – a cloud of pixie dust of sorts – whereby the components would effect electromagnetic fields and exchange photons amongst themselves to produce an effective capability equivalent to a monolithic system. This leap of the imagination is made somewhat easier by the observation that – distilled to their quintessence – most spacecraft missions involve little more than collection of photons emitted a source, some processing of this received signal, and subsequent re-radiation of photons to an interested target.

Globalstar). The difficulty of accurately assessing demand for satellite services is exacerbated by the increasingly lengthy development, manufacturing, and launch wait times, necessitating estimates far into the future. The accuracy of such projections are subject to market fluctuations, variability in the geopolitical climate, and technological innovation.

A fractionated architecture would solve this problem. The decision of how much capability should be designed into a spacecraft is not one that needs to be made years in advance of its launch. Instead, it is something that can be adjusted throughout the lifetime of the spacecraft by deploying additional modules. Thus, for instance, one could envision deploying an initial communications capability in the form of a power module, a T&C module, a C&DH module, and a handful of transponder payload modules. The decision to deploy additional transponder payloads could therefore be deferred until an initial operating capability is attained and the actual demand can be assessed. Individual transponder payload modules would allow for much finer "tuning" of on-orbit capacity to match demand than the double-or-nothing option available to the operator of a single large monolithic system. deWeck et al. have quantitatively shown that such incremental, scalable deployment can significantly impact the business case for commercial LEO communications systems. The result is generally applicable, however.

Likewise in the face of uncertainties other than user demand – for instance onorbit technical failure, component obsolescence, a reduction in program budget, a user requirement change, etc. – fractionated architectures offer the flexibility to adapt to the changed circumstance in real time. Traditional monolithic architectures, on the other hand, generally only allow for changes during the initial design phase. Thus, in response to each of the uncertainty factors faced by space systems enumerated *supra*, a fractionated architectures offers the postdesign option of substituting a module, augmenting the system with an additional module, removing a module from the system, or porting a module from one system to another.¹² These operations correspond to the various

⁹ Hardy, Q., "Surviving Iridium," Forbes Magazine (6 September 1999); Iridium LLC Chapter 11 Filing, Case No. 99-45005-CB, United States Bankruptcy Court, Southern District of New York (1999).

¹⁰ Their comparatively small size would also enable more rapid fabrication and responsive launch either aboard a small "tactical" launch vehicle, or by piggy-backing off any upcoming launch to an appropriate orbit. We address each of these topics in greater detail *infra*.

¹¹ deWeck, O., deNeufville, R., & Chaize, M., "Staged Deployment of Communications Satellite Constellations in Low Earth Orbit," *Journal of Aerospace Computing, Information, and Communication*, Vol. 1, pp. 119-136 (2004).

¹² Baldwin & Clark define six "modular operators" including the four we recite here – substitution, augmentation, exclusion, and porting – plus splitting (what we call fractionation), and inversion (analogous to changing the form of connectivity among the modules). They assert

manifestations of system flexibility: maintainability, scalability, and reconfigurability. Equivalent changes can, of course, be effected in a monolithic system as well, but they can only be made during the initial design of the system. Thus the key distinction between a fractionated and monolithic system is that the former retains elements of design flexibility throughout the operational lifetime of the system. This flexibility, in turn, provides robustness to the various uncertainties the system may encounter.

In order to make the various manifestations of flexibility meaningful system design parameters, a quantitative measure for comparing the flexibility of various designs is needed. Saleh, in his seminal dissertation and subsequent work,¹³ recognized that enabling the deferral of certain design decisions well into the operational life of a space system was analogous to acquiring and holding options¹⁴ which, when exercised, would produce an architectural change in the system (instead of triggering a purchase or sale of an underlying financial security). This notion of such "real" options has been recognized in the management world for some time, with the implication that a business plan which deferred a management decision into the future - when the uncertainty about the future state of the world is lower - was more valuable than a plan which required the decision to be made a priori. The option analogue is a particularly powerful one because there exists a relatively simple algorithm for quantifying the value of an option. The Black-Scholes formula provides an expression for the value of an option as a function of the price of the underlying asset, the volatility in the price of the underlying asset, and the exercise (strike) price of the option. So, for example, the option to augment a fractionated satellite with an additional payload module would have a specific value to its holder (the system operator, presumably) at any given time. 15 This value would be a function of the revenue which the additional module would generate if it were deployed (the price of the underlying asset), the future volatility in this revenue

that the operators are complete, i.e., a sequences of these operators can represent any design change to a modular system. Thus, they in effect formalize our assertion that fractionated architectures retain design flexibility throughout their entire development and operational life. Baldwin, C. & Clark, K., *Design Rules: The Power of Modularity*, MIT Press, Cambridge, MA, pp. 123-146 (2000).

¹³ Saleh, J., Weaving Time into System Architecture: New Perspectives on Flexibility, Spacecraft Design Lifetime, and On-Orbit Servicing, Doctoral Dissertation in Aeronautics & Astronautics, MIT, Cambridge, MA (2002); Saleh, J., Hastings, D., et al., "Flexibility and the Value of On-Orbit Servicing: New Customer-Centric Perspective," Journal of Spacecraft and Rockets, Vol. 40, No. 2, pp. 279-291 (2003).

¹⁴ Options meant here in the financial sense of the word. An option is a financial derivative – i.e., an instrument whose value is derived from some underlying security – which gives its owner the opportunity, but not the obligation to purchase or sell some other security (frequently shares of stock) at a fixed price.

¹⁵ It would be an American call option, i.e., an option to buy that can be exercised at any time prior to its maturity.

prediction (essentially the projected volatility of demand), and the cost to exercise the option and deploy the additional module (the strike price). If the revenue that the additional module would generate ever exceeds the cost to exercise the option and deploy the module, the option should be exercised. Prior to exercise, however, for the duration of its lifetime the option's value scales with the uncertainty in revenue (volatility of demand) – the greater the uncertainty, the more valuable the option.¹⁶

When a satellite operator invests into the development of a fractionated architecture he acquires not just the architecture, but also a bundle of options – which have real, computable pecuniary value – to make subsequent modifications to the architecture in response to uncertainty. The operator of a monolithic satellite with a comparable capability would not enjoy comparable options, nor be the beneficiary of their economic value. While the notion of "economic value" may seem anathema to operators of military systems, and while the absence of a functional market in which value can be determined from prices is a complicating factor, there exist methodologies for inferring equivalent revenue and value figures from communities which tend to think in terms of utility rather than dollar metrics.¹⁷

Portfolio Optimization - Diversification of Risk & Reliability Tailoring

The ancient adage "don't put all your eggs in one basket" has a firm theoretical underpinning. Consider the launch of a monolithic spacecraft and a fractionated spacecraft delivering the same capability. If each module of the fractionated spacecraft is launched separately – aboard a small responsive launch vehicle, or piggy-backed on other payloads – the cost impact to the program is smaller if

¹⁶ One of the excellent treatises on real options and their valuation, which explains the methodology and derives a variety of forms of the Black-Scholes formula, is Trigeorgis, L., Real Options: Managerial Flexibility and Strategy in Resource Allocation, MIT Press, Cambridge, MA (1996).

¹⁷ One approach to estimating the value of a particular system attribute or capability for which no market exists is as follows. First, a technique termed multi-attribute utility interviews is employed. The user is asked a series of questions eliciting their relative preference between different quantities of two attributes, A and B. Based on the responses, the relative elasticity between the user's preference for various quantities of A and B can be computed. The process can be repeated for attributes B and C, C and D, etc., until an attribute, say D, for which a market does exist is reached. Then based on the market valuation of D, a valuation of any other attribute can be computed through the relative preference elasticities. This technique is a hybrid of multi-attribute utility theory, see Thurston, D., "Multiattribute Utility Analysis in Design Management," *IEEE Transactions on Engineering Management*, Vol. 37, No. 4, pp. 296-301 (1990); Keeney, R. and Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*, Cambridge University Press, Cambridge, UK (1993), and contingent valuation methods, see Mitchell, R. & Carson, R., *Using Surveys to Value Public Goods: The Contingent Valuation Method*, Johns Hopkins Univ. Press, Baltimore, MD (1989).

one of the launches fails than if the single launch of the monolithic spacecraft were to fail. Stated another way, the variance of expected total launch costs for the fractionated spacecraft is lower than for the monolithic one. This results from the fact (or at least the supposition) that launch failures across multiple launch vehicles are independent random variables. One would suppose that a similar result would hold during orbital operations. And with respect to cost, it does. Thus, the on-orbit failure of a monolithic spacecraft necessitates the replacement cost (and associated launch cost) of the entire satellite. The failure of a single module of a fractionated spacecraft only leaves the operator with the replacement cost of the one module. A commensurate schedule effect due to the difference in replacement times for a large monolith versus a small component module can be expected. Thus the cost and schedule risk for a fractionated spacecraft over its entire lifetime is lower than for a monolithic one (assuming similar design reliabilities).

A far more interesting effect, however, is on spacecraft value. If a monolithic spacecraft suffers a fatal on-orbit failure it immediately ceases to deliver any value to its users (and hence to its operator). The failure of a single module in a fractionated system, however, leads to a less obvious result. In a purely heterogeneously fractionated system, it will similarly lead to a total loss of capability. In a purely homogeneously fractionated spacecraft, it will lead to a capability reduction proportional to the contribution of the failed module. In mixed systems, however, the covariance between the "capability contribution" of each module will dictate the extent to which the capability will suffer due to a particular module failure. As with option valuation, a financial analogue provides the requisite analytical framework. Per Markowitz, for an arbitrary number of investment vehicles with a given expected return, volatility of return, and covariance of return with each other investment vehicle, a set of optimal portfolios can be constructed.²⁰ Optimal portfolios are those which maximize total returns while minimizing the volatility of the total returns.

¹⁸ Brown considers this problem in more detail and finds that – without modeling the potential scenario where the fractionated spacecraft can still deliver some value even if one of its modules does not reach orbit – in order to get a 99.9% probability of a successful on-orbit operational capability, using reasonable launch cost and fractionation "mass penalty" assumptions, the expected launch costs are nearly a factor of two lower for the fractionated system than for the monolith. Brown, O., "Reducing Risk of Large Scale Space Systems Using a Modular Architecture," *Space Systems Engineering and Risk Management Symposium*, Manhattan Beach, CA (2004).

¹⁹ See Guikema, S. & Paté-Cornell, E., "Bayesian Analysis for Launch Vehicle Reliability," *Journal of Spacecraft and Rockets*, Vol. 41, No. 1, pp. 93-102 (2004).

²⁰ Markowitz, H., *Mean-Variance Analysis in Portfolio Choice and Capital Markets*, Frank J. Fabozzi Associates, New Hope, PA (1987).

Fractionated spacecraft provide a physical instantiation of such a "portfolio of modules," while also affording the opportunity to engage in portfolio optimization during design. The modules' functionality, variance in functionality, and covariance with every other module serve as perfect analogues for expected returns, volatility, and covariance of returns, respectively.²¹ Not only will the overall variance in system capability (i.e., value risk) for a such a "diversified" collection of modules be lower than for its monolithic counterpart, but it also gives the designer the flexibility to optimize the reliability of individual modules independently of the rest of the system. This is paradigmatically profound. Reliability design today is largely driven by a rather arbitrary specification of an exogenous requirement, or by attempts at rote maximization.²² If treated as a design parameter endogenous to the design process, the reliability of each module can be independently set to maximize the net value delivered by the system in light of considerations such as: the different paces of obsolescence for different technologies may make the early replacement of some modules (most notably C&DH) desirable; the cost of implementing a given degree of reliability may be starkly different across modules; and the degree to which the health of a particular module is vital to the capability of the overall system (i.e., its covariance with the other modules) may differ depending on the system architecture, degree of homogeneous fractionation, and types of connectivity.

One interesting approach to system diversification readily enabled by fractionated architectures is that of contractor diversity. It has been a long-standing precept of reliability engineering that components from different lot numbers and, if possible, different manufacturers should be used in the parallel strands of mission-critical redundant systems. This ensures that manufacturing defects that may affect a particular product lot are not repeated – in effect decorrelating the failure probabilities of the various redundant strands. Fractionation enables this application of this concept at the system level. The development and fabrication of the different modules of a fractionated system could easily be divided among multiple contractors without substantially affecting the complexity of the system integration effort, assuming a standardized inter-module interface is promulgated. This approach would

²¹ In addition to investment vehicles representing individual modules, the portfolio can also be thought of as containing the options (whose value is strongly correlated to the value of their underlying assets, i.e., the actual modules) from the preceding discussion on flexibility.

²² See, e.g., Saleh, J., Hassan, R., et al., "Impact of Subsystem Reliability on Satellite Revenue Generation and Present Value," *Journal of Spacecraft and Rockets*, Vol. 42, No. 6, pp. 1122-1129 (2005). For a more philosophical look at selecting the design reliability of engineering systems see Saleh, J., "Perspectives in Design: The Deacon's Masterpiece and the Hundred-Year Aircraft, Spacecraft, and Other Complex Engineering Systems," *Journal of Mechanical Design*, Vol. 127, No. 5, pp. 845-850 (2005).

ensure diversity of design and, presumably, diversity of design flaws. It could also have advantages for procurement competitiveness, as discussed *infra*.

There is little doubt that diversification is desirable, but optimal portfolio theory yields a set (a Pareto frontier) of portfolios with maximum aggregate returns and minimum volatility. The reliability and covariance optimization factors described above can aid in the detailed system design with optimal reliability for each module, but how does one initially assess the difference in value between systems with different degrees of diversification (i.e., degree of fractionation and extent of inter-module functionality covariance)? This question can be answered by analogizing the effects of system diversification to insurance. The objective of diversification is the reduction if cost and value variance, that is to say cost and value risk. Insurance is just a financial instrument with an effect analogous to the engineering "instrument" of fractionation. Thus historical insurance prices or a simple risk premium pricing model such as the CAPM²³ can yield a quantitative measure of the pecuniary benefit associated with a reduction in system value and cost risk due to diversification.

Spatial Distribution - Eliminating Fragility of Complex Systems

It is well known as an empirical matter (as well as a result from percolation theory) that a randomly-planted forest where the tree density is above a certain critical density is prone to catastrophic fires due to a random spark. Engineering design can make the forest, even with very high tree densities, robust to fires by intelligent placement of firebreaks – designed "structures" which in this case are gaps cut through the forest. Flexibility to add firebreaks throughout the life of the forest or even as a fire is propagating – rather than attempting to design them in during the initial planting – can enhance robustness further by delaying firebreak placement decisions to a point in time when the uncertainty as to ultimate tree size, prevailing winds, and other relevant environmental factors is lower. Diversification of tree types can enhance robustness even further if the various species burn at different rates or have different susceptibility to combustion. These measures, in effect structures added to the forest through engineering design, can greatly enhance its robustness to a known risk – sparks. At the same time, this additional complexity introduces new sources of fragility which did not previously exist - sensitivity to imperfections in the firebreaks, for instance, or the probability distribution of sparks, or humidity variations which may affect fire propagation speed.

²³ The Capital Asset Pricing Model (CAPM) assumes a simple linear relationship between risk and the associated market premium. The line is typically calibrated with two points: the return on (virtually) risk-free Treasury bills and the so-called market portfolio (which is an aggregation of all the securities in a particular market). See, e.g., Meyers, S. & Brealey, R., *Principles of Corporate Finance*, McGraw Hill, pp. 194-203 (7th ed., 2003).

We previously defined fragility as the susceptibility of complex systems – which may be robust to known uncertainties – to unanticipated (sometimes termed "emergent") failure modes, frequently arising either from cascading failures or unmodeled component interactions. Although rigorous modeling, simulation, and integrated system testing can occasionally identify fragilities, even ameliorating actions have the potential to introduce new fragilities into the system. The only sure way to reduce the fragility of the forest is to reduce the density of the trees. Likewise, the only certain way to reduce the fragility of a complex system such a spacecraft is to make it less complex. Is it possible to do so in any meaningful way without sacrificing capability?

To answer this question we must necessarily digress briefly to define "complexity." Complexity is not a mere metric for the number of component parts of a system. If it were, a heterogeneous material, e.g., a block of metal, which contains an abundance of molecular components, would be considered immensely complex. Instead, "[i]t is the extreme heterogeneity of the parts and their organization into intricate and highly structured networks, with hierarchies and multiple scales"24 that constitute the essence of complexity. Thus, discounting financial markets and human organizations from the space of engineering systems, let us assume arguendo that the most complex system ever developed by mankind is the Internet in its entirety. A distant second might be the Space Shuttle. Yet the Shuttle is far more prone to catastrophic failures resulting from fragilities than is the Internet.²⁵ Why? One can imagine that if the aggregate complexity of the Internet were to be replicated in a single, tightly integrated supercomputer, it too would be a highly fragile system. Instead, the Internet reduces its fragility through spatial distribution of complexity. Analogously to the tree density in our forest fire example, the "complexity density" of the Internet is much lower than of the Space Shuttle. Naturally, the physics of hypersonic flight do not lend the latter the luxury of a distributed architecture, but the comparatively docile environment of space does afford satellites the opportunity to distribute their complexity and reduce mechanical, thermal, and other undesirable component interactions while preserving the desirable exchange of data, power, and certain forces and torques. Thus, with the implementation of an additional layer of abstraction - the protocol for interfacing

²⁴ Carlson, J. & Doyle, J., "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 1, p. 2539 (2002).

²⁵ It is noteworthy that the Internet does experience occasional catastrophic failures due to fragility that arises from the interface between its myriad of nodes. In distributing the complexity of a monolithic system across numerous nodes or modules, some complexity – and hence new fragilities – must be added in the interfaces and protocols. See, e.g., Doyle, J., Alderson, D., et al., "The 'Robust Yet Fragile' Nature of the Internet," *Proceedings of the National Academy of Sciences*, Vol. 102, No. 41, pp. 14497-14502 (2005). Therefore, when distributing a satellite one can imagine fragilities arising from environmental interactions with the crosslinks, for instance.

the modules to one another – the distributed nature of the fractionated spacecraft can reduce its fragility and consequent susceptibility to catastrophic failures. In effect, our metaphorical forest can be thinned without reducing the number of trees simply by spreading it out.

A more straightforward, but nonetheless important effect of distribution is the physical decoupling of the payload module or modules from the rest of the spacecraft and from each other. Since payload needs tend to drive the pointing accuracy requirements in a monolithic spacecraft, they can be relaxed for the non-payload modules of a fractionated one. This can significantly simplify and shrink the attitude determination and control system (ADCS) for most of the spacecraft modules. And even for the sensitive payload the problem is easier – as the smaller inertias of the module alone lend themselves to faster control loops, and the coupled structural dynamics and vibration problems endemic to large spacecraft are alleviated.

This physical decoupling also enables the development of classified payloads separately from the rest of the spacecraft – without imposing associated security costs and constraints on the modules (including other payloads) that don't require them. The classified payload can be launched separately from the rest of the spacecraft, and remain isolated in orbit except through the standard communication, power, and whatever interfaces are necessary. There is little technological challenge to making the interfaces secure and spoof-proof.

Cost Considerations

Our discussion hitherto has focused on the value of fractionated systems in contrast to their monolithic counterparts. What about the cost? The necessary replication of certain subsystems (e.g., thermal control), the elective replication of others, and inter-module interface hardware even across a fully heterogeneous system would lead one to expect the total mass of the fractionated spacecraft to be higher than a comparable monolithic one. Interestingly, preliminary design studies have shown that this may not always be the case. The weight savings from a reduction in ADCS system mass (primarily flywheel mass) resultant from separating the payload from the rest of the spacecraft, particularly in certain classes of earth- and deep space- observing spacecraft, can be very significant. In any event, however, the assumption that spacecraft cost scales roughly linearly with its mass is an artifact of the *status quo* in acquisition, design, and manufacturing processes.

Fractionated space architectures have the potential to commoditize the space industry. As we discussed, flexibility and diversification are likely to make fractionated systems more tolerant of component- or module-level failures

without compromising overall mission reliability. Their distributed nature and standardized interfaces will make similar modules easily portable across many different spacecraft. Thus, a universal power, C&DH, or T&C module could be offered; users requiring more capability or higher reliability would simply deploy multiple such modules to work in parallel or as redundant backups. The combination of reduced reliability requirements and the demand for a multitude of similar or identical products are a magical combination enabling the utilization of mass production techniques prevalent in other high-technology industries such as automotive and computer. The limited application of some mass production techniques was met with considerable success in Lockheed Martin's construction of the 79+ Iridium satellites. Even without fundamental process changes, learning curve slopes of 90% to 95% are typical in the space industry. One might reasonably expect them to drift towards the 75% to 85% range, more typical of the aircraft and mass market high-tech industries. Analogously to a Beowulf cluster²⁶ delivering the capability of a massive supercomputer - while undoubtedly replicating some of the overhead components such as cooling fans and unnecessary input/output capability - at a tiny fraction of the cost, the commoditization and mass production of spacecraft modules may well offset whatever cost penalties are incurred in replicating some of the infrastructure enabling of fractionation.

Although we focus our attention in this paper on space architectures, the launch infrastructure is an inextricable and costly component of the lifecycle of the satellite system. We discuss the effect that fractionated architectures may have on the launch market in the next section. Here we focus on the reverse relationship the contribution of launch cost to the cost of fractionated spacecraft. There are three primary means of placing fractionated spacecraft into orbit. First, they can be launched simultaneously aboard a single large launch vehicle - analogously to a monolithic spacecraft. This approach fails to capitalize off the diversification advantage due to the decorrelation of potential launch failure events for individual modules. If the aggregate mass of the fractionated system exceeds that of the monolithic one, it will incur a launch cost penalty using this approach. Second, individual modules of the fractionated system can be launched as "piggy-back" payloads alongside existing launch opportunities. The EELV Secondary Payload Adapter (ESPA) ring is one such approach to integrating small satellites as secondary payloads alongside existing Delta IV or Atlas V payloads.²⁷ Dispersing the modules of the fractionated satellite in this manner preserves the launch diversification advantages, but does not exploit the full

²⁶ A popular open-source architecture for effecting massive parallel processing capability with cluster of personal computers. See, e.g., http://en.wikipedia.org/wiki/Beowulf_cluster for a helpful description.

²⁷ Anonymous, Secondary Payload Planner's Guide for Use on the EELV Secondary Payload Adapter, DoD Space Test Program (2001).

responsiveness potential of the fractionated architecture, and is contingent on the availability of launch opportunities into the vicinity of the desired orbit. The third approach to launching fractionated satellites is by placing individual (or a small number, depending on size) modules into orbit with dedicated small satellite launch systems, a number of which are presently under development in response to the Air Force's *Mission Need Statement* for *Operationally Responsive Spacelift*. Although the cost projections for these launch vehicles are promising, none have hitherto entered operational service. Fractionated spacecraft are likely to enjoy a symbiotic relationship with small responsive launch systems – the former may provide the launch volume to make the latter economically competitive, and the latter may enable the full responsiveness and flexibility of the former.

As a final note on the subject of cost, we observed previously that one approach to achieving diversification of design approaches and component choices was to spread the development of the various modules of the fractionated spacecraft across multiple contractors. One of the additional benefits of such an acquisition approach is increased competition among the contractors, not just because of the increase in the number of competitive opportunities, but also due to the reduction in the barrier to market entry which a smaller firm or one from outside the space industry must incur to successfully compete against the major incumbents. Such a highly competitive marketplace – unlike the oligopoly of today – should effect substantial downward cost pressures on its participants, both in terms of reducing profit margins, improving efficiency, and stimulating process and technology innovation.

New Paradigms

Although the focus of this paper is primarily to make the case why an acquisition official, a technology strategist, or a satellite designer might select a fractionated architecture versus a monolithic one, we are mindful of the broader ramifications if spacecraft fractionation were to gain widespread acceptance. From an acquisition perspective, for instance, fractionated satellites would be enabling of very rapid design-build-test development cycles. This would, in turn, permit a radically different approach to acquisition of a variety of space systems. Whereas the notion of spiral acquisition has become widespread as a means of making acquisition processes robust to technology development and other uncertainties, most major programs manage to effect at most two or three top-level spirals. In reality, the structure of the acquisition program can be optimized given an estimate of the risks and the costs associated with the development of multiple

²⁸ Kolodziejski, P., "Operationally Responsive Spacelift for the U.S. Air Force," AIAA-LA/SSTC-2003-1002, 1st AIAA Responsive Space Conference, Redondo Beach, CA (2003).

product blocks.²⁹ The rapid design-build-test cycles of fractionated systems would enable highly-spiral acquisition if warranted by the particular risk environment and uncertainties associated with a given program. Another innovation in acquisition that may accompany fractionated systems is incremental deployment of capability. Not only is quick design-build-test possible, but so is rapid design-build-fly. Thus, launch costs permitting, early blocks of particular modules could be deployed for an extremely responsive initial operating capability, only to be replaced by subsequent, more mature versions.

Fractionation can be an enabler for responsive space. Conventional wisdom suggests that responsiveness is the ability to quickly develop and launch orbital payloads. We dissent from this narrow formulation. Although shortening the development and launch timelines is one instantiation of the solution, we understand the problem of responsive space to be one of effecting a space-based capability rapidly in response to uncertainty. This broadens the solution space; it permits us to consider alternate - and undoubtedly complementary - means of enabling responsiveness across a wide range of systems, large and small. The conventional view of responsive space is predicated on one manifestation of uncertainty, that of a combatant commander faced with a temporally and geographically uncertain threat. Hence, resultant efforts focus on the need for launch on demand in response to a particular tactical threat, at a specific location, on short notice. But there are other manifestations of uncertainty which vex not only the warfighter in the field, but also the acquisition official, the spacecraft designer, the manufacturer, the tester, and the operator. Uncertainty exists throughout the entire lifecycle of a space system, and therefore the need for responsiveness is omnipresent from its cradle to its grave. Fractionated spacecraft offer an architectural, rather than process-based, solution to the problem of responsiveness - not just by shrinking spacecraft development timelines and enabling launch with small tactical vehicles, but also by making the spacecraft architecture fundamentally quicker at adapting to uncertainty.³⁰

Furthermore, as we mentioned previously, fractionated architectures could be enabler that makes responsive space launch economically viable by providing requisite launch rates and volumes to produce meaningful learning curve effects in the launch vehicle supply chain, and to amortize the costs of development and

²⁹ In fact, each spiral can be thought of as creating a real option to continue product development or refinement (potentially with a new set of requirement or budget), terminating the program altogether, or delivering the product as is. Each spiral therefore adds value, but incurs the additional cost of going through another design-build-test cycle.

³⁰ These ideas are elaborated at greater length in Brown, O. & Eremenko, P., "Fractionated Space Architectures: A Vision for Response Space," AIAA-RS4-2006-1002, AIAA 4th Responsive Space Conference, Los Angeles, CA (2006).

infrastructure. Fractionated spacecraft could be the *raison d'être* for operationally responsive spacelift. While the utility of tactical responsive payloads is inherently limited by their small size, fractionation removes this constraint. Thus, a theater commander could deploy a payload or multiple payload modules on short notice with responsive spacelift, only to have the payload inserted into an existing on-orbit infrastructure supplying it with power, telemetry, computational and data storage resources, and perhaps even stationkeeping. Taking this notion to its logical extremum, one could envision a global infrastructure of spacecraft "utility" modules, supplying all the necessary support functions to enable a free-flying payload module to be inserted into this network on an ad-hoc basis – anytime, anywhere.

And finally, but perhaps most importantly as far as new operational paradigms are concerned, fractionation enables the construction of space systems well in excess of what can presently be put in orbit with the largest launch vehicles in the arsenal. It is, in effect, a lower-risk and more flexible replacement for on-orbit manufacturing or assembly of super-satellites for mission applications which may not necessitate large, rigid structures.³¹

IMPLEMENTATION

Value-Centric Acquisition

From our preceding discussion, it is obvious that the compelling proposition for fractionated spacecraft architecture lies in the enhanced value – stemming from the flexibility and robustness – that they offer throughout their operational life. In a functional competitive market, product investment decisions are based on the net value (lifecycle value minus cost)³² that a product delivers over its lifetime. The existing DoD acquisition framework, however, does not attempt to mimic the behavior of a functional market and reward value maximization; instead, it is aimed at minimizing cost for a given set of requirements. If the requirement process was analogous (and appropriately weighted) to the value-seeking preferences of a rational operator of satellite services, the result would be identical. The requirements formulation process, however, is neither driven by

³¹ Thus, for instance, fractionated architectures are notably unsuitable for manned spacecraft applications where a large pressure vessel must be constructed. Their applicability to interplanetary missions is also questionable due to the dominance of the propulsion requirement across the entire system.

³² More precisely, net value is the net present value (NPV) of the value streams less the cost streams. The value streams include the revenues generated by the system, plus the option value and a value premium if the volatility of the value stream is low (i.e., value risk insurance). The cost streams include all the actual expenditures needed to develop, construct, launch, and operate the system, including the strike price of any options which were exercised, plus a cost premium if the volatility of the cost stream is high (i.e., cost risk insurance).

satellite operators, nor is it intended to span system *attributes* (as opposed to system *capabilities*). Instead, the requirements process is conducted with the users' interests in mind. While the user may the ultimate consumer of satellite services, he is in no position to formulate requirements for system attributes such as maintainability, reliability, flexibility, etc., but only for system capabilities such as bandwidth, resolution, and revisit rate. The result is that between two systems which deliver the identical required capability, the DoD would choose the less expensive one over the one that delivers the most net lifecycle value, even if both fall within the allotted program budget, though it is clearly a suboptimal choice from a value-centric perspective.

Interestingly, this acquisition mindset focused on minimizing cost for a given capability requirement appears to have carried over from military to commercial aerospace. Perhaps this is due to institutional inertia of the major contractors, or maybe due to a shared cast of characters. The prolific use of flawed design metrics such as satellite cost per operational day and cost per transponder³³ is a manifestation of the cost-centric mindset and leads to decisionmaking which is clearly not tailored to maximizing return on investment and hence shareholder value.

As we noted *supra*, the cost proposition for fractionated systems is ambiguous – that is to say there is undoubtedly some duplication of hardware that enables the wireless interaction of the modules, along numerous countervailing trends the net effect of which is nearly impossible to assess *a priori*. The value proposition of fractionated spacecraft, however, is unambiguous. But in order for it to be recognized, acquisition decisions must rest on a net value criterion. That is the first enabler to the implementation of fractionated architectures.³⁴

³³ Saleh, J., "Flawed metrics*: satellite cost per transponder and cost per operational day (*for guiding design decisions)," *IEEE Transactions on Aerospace and Electronic Systems* (accepted), Vol. --, No. -- (2006).

³⁴ We have elsewhere attempted to construct a rudimentary quantitative net value model for fractionated communications satellite. We found non-trivial amounts of lifecycle value emanating from the non-traditional sources such as flexibility, reduced variance of cost and revenue streams, and incremental deployment. With an admittedly simplistic parametric cost model, we found that even for the most basic fractionated architectures easily yielded value increments vice their monolithic counterparts that offset the non-recurring and recurring cost penalties. Our notional model is documented in Brown, O., Eremenko, P., & Roberts, C., "Cost-Benefit Analysis of a Notional Fractionated Space Architecture," AIAA-2006-5328, 24th AIAA International Communications Satellite Systems Conference, San Diego, CA (2006). See also Mathieu, C. & Weigel, A., "Assessing the Flexibility Provided by Fractionated Spacecraft," AIAA-2005-6700, AIAA Space 2005, Long Beach, CA (2005).

Cluster Flying

The other enablers of fractionation are technological. We use the term cluster flying to refer to persistently proximate orbital positioning of multiple satellite modules in passively stable, Keplerian orbits. Such orbits can be constructed by effecting a small perturbation to modules which are otherwise in co-altitude orbits. A small eccentricity change can create a co-orbiting cluster in the plane created by radial and in-track relative motion, while an eccentricity perturbation would create motion in the radial and cross-track plane.³⁵ Such orbits are called halo orbits and permit a cluster of arbitrary size to be stationkept with only second-order ΔV expenditures to compensate for drag, differential force, oblateness, and third-body effects.

We are careful to distinguish between our notion of cluster flying and the more commonly discussed concept of formation flying. As will be readily apparent from the subsequent discussion, fractionated architectures do not generally require precise maintenance of relative module attitude or position, but only their determination with sufficient accuracy to enable pointing of power transmission links. Thus, relative drift of the modules due to higher-order orbital disturbances is perfectly acceptable so long as relative distances and orientations do not exceed the ranges supportable by the cross-links, and so long as collision avoidance can be ensured. This alleviates the technical challenges of the relative stationkeeping problem, and instead simplifies to a rather moderately difficult question of relative navigation.

Data Transmission

Data exchange is the most basic form of connectivity among the modules of a fractionated spacecraft. It simply involves replacing the data bus of the monolithic spacecraft with a series of wireless data links among the several modules of the fractionated one. Data exchange alone would permit the heterogeneous fractionation of at least T&C, C&DH, and payload modules. Not unlike the inter-satellite crosslinks used by Iridium and TDRSS, data exchange poses little in the way of technological challenge. A variety of technologies fall within the tradespace, including low-power, omnidirectional, spread-spectrum links analogous to IEEE 802.11 (which relax much of the relative orientation requirement between modules and permit ad-hoc addition and removal of nodes), or the emerging ultra-wideband (UWB) technology (which can also provide centimeter-precision relative position information between transceivers). Alternatively, if power is also exchanged between modules (see *infra*), the

³⁵ Wertz, J.R., Mission Geometry; Orbit and Constellation Design and Management, pp. 519-525, Microcosm Press (2001); Bekey, I., Advanced Space System Concepts and Technologies, pp. 19-20, The Aerospace Press (2003).

communications signal can be modulated on top of the optical or RF power beam. The overriding concern in exchanging data wirelessly is the effective exposure of the link to external observation. While encryption can ensure that the information content of the signal is not intercepted, care must also be taken in its design to provide interference-resistance capabilities. Notably, both spread-spectrum and ultra-wideband technologies provide excellent signal robustness.

Fractionated Navigation

Navigation functionality can be fractionated into a separate module responsible for determining its position and attitude in an absolute (inertial) reference frame. It could be the only module aware of its inertial position, with the rest of the fractionated spacecraft determining its position and attitude relative to the navigation module. Synergies with data, power, and perhaps even force/torque crosslinks could be exploited to yield relative distance and orientation information with minimum additional hardware.

Distributed Computation & Data Resources

While relegating intensive computation and data storage to a specialized C&DH module (or a plurality thereof) may be the most effective means of capitalizing off Moore's Law in fractionated system, a fleet of fractionated satellite modules is nonetheless aptly viewed as a set of distributed agents that together perform the processing functions of a conventional monolithic spacecraft. The seamless utilization of resources across these agents to effect a real-time processing capability that is adaptive to resource fluctuations and that provides a secure (from both a cryptographic and authentication perspectives) application environment is a key enabler for the fractionation paradigm. One prospectively appealing instantiation of such a distributed environment would transparently incorporate ground nodes as a natural extension of the "virtual spacecraft," and – subject to potentially increased latency and reduced availablity of data links – allocate appropriate tasking to these cheaper and more abundant resources.

Power Transmission

Fractionation of power generation capability into a separate module requires its wireless dissemination throughout the rest of the virtual spacecraft. A variety of means for wireless power transmission are in the trade space, the choice among which is driven largely by inter-module distance. Preliminary studies appear to favor radio frequency transmission at distances below several hundred meters, with V-band, W-band, or higher frequencies being especially appealing from a rectenna size perspective. Beyond inter-module distances of several kilometers, laser transmission to a tuned band gap photovoltaic cell appears preferable. If

power transmission only during sunlight hours is acceptable, then solar collection, concentration, and reflection (i.e., without first converting the solar energy to electricity aboard the power module) to a heat engine receiver promises considerable efficiency improvements over both RF and laser links. Induction offers yet another option which, while conveniently omnidirectional, is only efficient at extremely close ranges where the non-radiating component of the time-varying magnetic field is dominant.

Force & Torque Transmission

Currently on the technological horizon is the fractionation of propulsion. Remote forces and torques can be effected from a designated propulsion module to the rest of the fractionated cluster. A viable approach appears to be electromagnetics, as demonstrated in the laboratory by Miller et al.³⁶ Each module is equipped with three orthogonal electromagnetic coils which, when energized, can create an effective magnetic dipole in arbitrary orientation. The interaction of a pair of such dipoles produce torques and moments (which can be reacted with a reaction wheel if the desire is to induce motion in only one of the modules) that can be used for stationkeeping or cluster reconfiguration purposes. Another approach for inter-module force transmission relying on electrostatic forces has been proposed for use in GEO by Parker et al.³⁷

Demonstration Program

The Defense Advanced Research Projects Agency (DARPA) has been studying the fractionated architecture concept and is poised to commence an initiative entitled System F6 – short for Future Fast, Flexible, Fractionated, Free-Flying Spacecraft united by Information eXchange, and incidentally a tornado of unimaginable strength on the Fujitsu scale³⁸ – that will mature the associated technological, architectural, and organizational advancements necessary for an on-orbit demonstration of a fractionated spacecraft. F6 will explore a rapid, multi-spiral design-build-test program structure, and will require the utilization of explicit quantitative system value models to support design decisions. We anticipate the formal start of the System F6 program at the beginning of FY2007 to culminate in an orbital demonstration in the FY2008 – FY2009 time frame. The

³⁶ Kong, E., Kwon, D., et al., "Electromagnetic Formation Flight for Multisatellite Arrays," *Journal of Spacecraft and Rockets*, Vol. 41, No. 4, pp. 659-666 (2004).

³⁷ Parker, G., King, L., & Schaub, H., "Charge Determination for Specified Shape Coulomb Force Virtual Structures," AIAA-2006-1891, 47th AIAA Structures, Structural Dynamics, and Materials Conference, Newport, RI (2006).

³⁸ F6 is also the fictitious mountain which reveals to its climbers more about themselves than they ever knew (or, perhaps, ever cared to know) in the brilliant play by Auden, W.H. & Isherwood, C., *The Ascent of F6: A Tragedy in Two Acts*, Faber & Faber (1937).

end goal of the F6 program is to fabricate and space test a microsatellite-scale fractionated space system.

Conclusion

The advent of the integrated circuit some four decades ago set off mankind's insatiable thirst for computational power. The quest to quench this desire led to the development of increasingly more sophisticated computers. Microchips sprouted ever greater numbers of transistors, choking buses, and forcing memory banks to struggle to keep up. The novelty of micro- and minicomputers was quickly trumped by the sheer computational prowess of supercomputers. And so the trend continued. In a matter of two decades, however, this drive towards greater processing power culminated in mammoth mainframes whose rapidly increasing complexity, fragility, and cost quickly outpaced the capability gains. A scant few years into the second decade of the era of the integrated circuit, the availability of inexpensive, mass-produced microcomputers, and the advent of fast, seamless internetworking ensured the relegation of the large monolithic mainframes to obsolescence and obscurity. Spacecraft have followed a trajectory that is uncannily parallel (and, of course, technologically intertwined) to the history of high-end computing. Borrowing the historical analogy, we posit that the era of distributed space architectures has likewise arrived. The gargantuan monolithic systems deployed to orbit today have grown too large, too complex, too fragile, and consequently much too expensive; furthermore, these trends have not been offset by commensurately rapid growth in capability. On such systems the proposition of responsive space cannot viably rest; we instead tender the paradigm of fractionation as the answer.

Acknowledgements

The authors would like to acknowledge the intellectual contributions of the following individuals to the ideas contained herein: Joseph Saleh, Ivan Bekey, Lyle Abramowitz, Talbot Jaeger, Lisa Hill, Annalisa Weigel, Charlotte Mathieu, Christopher Roberts, Roger Hall, Richard Golding, Theodore Wong, David Miller, and Ray Sedwick.